ICANN73 | Virtual Community Forum – Tech Day (2 of 3)
Monday, March 7, 2022 – 10:30 to 12:00 AST

KATHY SCHNITT: Hi, everyone. Thank you for joining Part 2 of Tech Day at ICANN73. As a reminder, this session is recorded and follows the ICANN expected standards of behavior. Participants are welcome to post their questions or comments using the Q&A pod of via Zoom chat. Again, thanks for joining, and I'll hand the call back over to Dr. Eberhard Lisse.

EBERHARD LISSE: Thank you. Welcome back. Without further ado, Andrew McConachie has the floor. You can share your screen or use the screen provided by Kim.

ANDREW MCCONACHIE: Sure. One moment.

EBERHARD LISSE: We can hear you as well.

ANDREW MCCONACHIE: That's important. Let me get a bit of video going. Do you see video now?

EBERHARD LISSE: We see video and your shared screen.

ANDREW MCCONACHIE: Okay, perfect. My name is Andrew McConachie. I work for the ICANN Policy Department where I support the Root Server System Advisory Committee and the Security and Stability Advisory Committee. I'll be talking about a new website, rssac002.root-servers.org, which hosts some visualizations of RSSAC002 data.

First, I'm going to talk about what RSSAC002 data is, and then I'll talk about the software I wrote, and then finally I'll do kind of a little case study on Chromium queries to the root and what we can see of them in RSSAC002 data.

Okay, so RSSAC002 is a specification produced by the RSSAC for data that root server operators publish on their own. It's collected daily. It's collected over a 24-hour period, and it's primarily a bunch of counts. So it has the granularity of the day.

There's nothing to do with names. It's really just counts of queries and responses, counts of different kinds of RCODEs responded with, unique IP sources seen, these kinds of things.

It's currently on Version 4. And it is an active specification, so it is actively maintained. There's a nice repository of all this RSSAC002 data there in GitHub maintained by the RSSAC Caucus.

What I did, this basically consists of two parts. There's a web API and then there's a website hosting charts that use the API. It's all open source, and I wanted to include all these people's names who gave me a bunch of really great feedback. I've been developing this for months now, and it was really useful to get a lot of really good feedback from

people on how to make this data useful in a visual sense. So these people here really helped me with that.

And I'm still interested in feedback. Feel free to go to the website. Play around with the charts. If you find any bugs, please raise an issue on GitHub or contact me directly at that email address. I'm always interested in people's feedback.

So the web API is an HTTPS interface to the RSSAC002 data. It returns JSON formatted data. The RSSAC002 data is YAML. That's how it's published. I read it in every day and then chew on it and then produce what's essentially serialized PHP data structures. And then when the requests come in over the web, it returns JSON.

I made a decision to start January 1, 2017. I wanted to start. You know, all this data is time-series data, and it's published every day. So you have to start somewhere. I think the first version of RSSAC002 came out in 2014, but it took a while to get implemented and whatnot. But that gives us a good five years to play with right now. So we have five years of data right now, and going forward we'll continue to produce data.

And again, the source code and documentation are on GitHub. And this is really independent from the charts. So if you don't like my charts for whatever reason or you just want to make your own, you can send calls to this API and get RSSAC002 data in a JSON format if that's your thing.

So a little bit about the charts. These are what's hosted on the website, what you see if you go to the website, just index.html. They of course pull data from the RSSAC002 web API. And as I said, we started in 2017

January, and it goes up to 21 days before today. So there's always three weeks of lag there, and that's just to make sure that we have good data. So there's always a little bit of lag there.

Everything is manipulable in the browser. Lots of time-series. Some exceptions. There are some pie charts and whatnot which I'll show you, but it's mostly time-series. And it's done entirely in JavaScript.

Okay, so that's a bit about RSSAC002 data, what it is, and then also the software that makes the charts. Now I thought it would be fun to just look at Chromium queries to the root. Not because I'm particularly interested in Chromium queries, and there are other people who have done better investigations of these queries than I'm going to do. I'm not really interested in why these queries are being generated. I'm interested in what we can see of these of these queries in RSSAC002 data and how we can visualize that.

So for a bit of background, Chromium is the open source engine for Google Chrome and Microsoft Edge. It used to generate DNS queries for random strings between 7 and 15 characters which was implemented to detect if the browser was behind NXDOMAIN hijacking.

There was a change introduce November 4, 2020 to no longer send those queries. And you can see in the data that that change happened pretty quickly, and we'll see that in the charts.

So it was a bit of a research question: What can we see of the Chromium queries in RSSAC002 data? As I said earlier, RSSAC002 data is very

coarse. It's basically just a bunch of increment encounters. But we should be able to spot big systemic trends.

And it's rather telling that the fix to Chromium rolled out in on November 4, 2020, and then traffic to the root server system peaked on November 4, 2020. With the exception of some DDoS attacks which happened prior to November 4, 2020, which increased past that. That's why that asterisk is there. But essentially we see traffic to the root server system peak on November 4, 2020, right when the fix to Chromium is rolled out.

I chose two 12-month periods for comparison. The first is rather obvious, November 4, 2019, to November 4, 2020. Then I figured wait three months to be certain that the fix is rolled out, and then take another year from February 1, 2021, to February 1, 2022.

This is just total traffic to the root server system, and it's by week. We have daily granularity, but that really makes a pretty spiky chart. So you can smooth it out by just looking at weeks as data points. You can see there's that peak there, and that's right when the change to Chromium was released. Then you see a sharp drop off, and then you see this continual gradual drop off.

Here's the same chart again but with each different root server identifier on it for a different color and they're stacked. One of the things about these charts, I see these charts a fair bit and they're a little bit…they kind of give you the impression that the line on the bottom didn't experience much of a drop. But they're a little bit misleading in that regard because actually if you look at each RSI, each letter

individually, you'll see that they all experienced a bit of a drop. But you see it most at the top of the stack because there's this compounding of the peaks that works its way up the stack.

These are all available at the website, and you can explore these further if you wish.

This is basically a comparison of NXDOMAINs and NOERRORs. You see that in the two different periods we've chosen to look at there's a 13% drop in NXDOMAIN responses. So we can probably assume that the Chromium query has a lot to do with that.

What we're looking at here is IPv4 queries received over unique sources. So basically, how many queries do we get per source, per unique IPv4 source. It has dropped. It doesn't look like much because it's a logarithmic scale on the Y. But in some, it looks like it has dropped by about two-thirds. I think F is on the top there. We can see they start off 2020 at about 22,000 per source, and then toward the beginning of 2022 they're somewhere around 8,000 per source.

Really the only thing you can really get from packet sizes, or at least the only thing I've really been able to [DANE] from looking at packet size data in the RSSAC002 data is an increase in diversity. So if we look at our first period, we see that we have about 56% of UDP queries coming in between the 32 and 47. Then we look at our second period and it has gone down by about 11% there. And then that has been spread across to these other sizes.

Again, with RSSAC002 data it's hard to really say much more than…you can't do interesting correlatives like this percentage of packets in this size produced NXDOMAINs or something like that. You're just looking at individual counters which can't be directly correlated with one another so easily. And then this is the UDP response sizes as well. So we just see more diversity at the high end after the Chromium queries stopped.

And that's it. This is really just meant to kind of introduce the website and introduce some of the things you can do. I'm happy to answer any questions people might have.

EBERHARD LISSE: Thank you very much. I was looking for my mouse pointer to find the unmute button. Interesting stuff. If you could use JSON, it's relatively simple to run. It's for other things to make rough like R, the statistical language which can import JSON very nicely which I use for my own stuff a little bit.

If you go two slides back, there was a huge spike on one name. This one. What spike is this? This one name server, what spike is this?

ANDREW MCCONACHIE: I'm going to guess that's probably a DoS attack of some sort. I don't know. But, yeah, lots of queries from same number of sources.

EBERHARD LISSE: Okay. There are two hands raised. I have asked the first one to put the question in the Q&A pod, which hasn't happened. So that's a problem.

We will only take questions from the Q&A pod, so please post them there. Okay, I don't see anything there, so we'll give it a few more seconds. We're not in a hurry. There we go. Ken Renard: "During your work, did you discover any matrix that might be useful to add to RSSAC002 data?" I'm just reading this aloud so that this goes into the record.

ANDREW MCCONACHIE: That's a good question, Ken. Nothing really comes up off the top of my head. Again, I think that it's important that these do remain simple counters. There's always a desire to have more data. As engineers we just want more and more data, but you don't want to be too burdensome to root server operators. So I'd really have to think about that. But, I mean, the RSSAC Caucus does review RSSAC002 stuff every two years. Not that they necessarily produce a new version, but it is reviewed in case there's a new version needing to be created. So I'll have a think about that. The next time we review it, I'm sure we'll talk.

EBERHARD LISSE: There is a question in the chat which I find intriguing. "Can this tool work for any name server logs?" I think that's a misunderstanding.

ANDREW MCCONACHIE: No.

EBERHARD LISSE: I think there is a misunderstanding because it's not a name server log that is analyzed. It's data that the root server operators provide to you, isn't it?

ANDREW MCCONACHIE: It's data that the root server operators provide publicly, and it has a unique specification. I mean, it is just YAML. So if you have a library that can parse YAML, you can parse this data. But it's not intended for any other authoritative server operator really.

EBERHARD LISSE: I'm quite sure there are a number of tools flying around open source and so on that can analyze name servers. Even I did this a few years ago, some simple stuff. That's not a problem. But the question was interesting because it's a misunderstanding. This is data that is provided by the root server operators in terms of an advisory from the stability committee. This website then makes it palatable, I think.

ANDREW MCCONACHIE: That is interesting.

EBERHARD LISSE: Anyway, yes, sure. I just wanted to make sure that this particular individual has the answer. Any other questions? I would then, if not, thank you very much and ask Craig Schwartz to unmute himself and take the floor.

CRAIG SCHWARTZ: Excellent. Thank you. Hello, everyone.

EBERHARD LISSE: Turn your video on, please, if you can.

CRAIG SCHWARTZ: Sorry about that. There we go. Am I going to be driving this stack, or are you?

KATHY SCHNITT: Craig, I have it queued up for you.

CRAIG SCHWARTZ: Okay. Thanks, everyone, for the opportunity to speak with you today and for the introduction provided a little bit earlier. My name is Craig Schwartz. I'm the managing director of fTLD Registry Services, and we manage the .bank and .insurance top-level domains.

I'm here to talk with you today about DMARC for public suffix domains, which is a new type of security that we've been working on for about the last four years and when implemented should mitigate the most pervasive form of DNS abuse which is phishing.

That financial institutions continue to be the most targeted sector online for abuse is one of the reasons that a coalition of banks, insurance companies, and financial services trade associations came

together and formed fTLD in 2011 and ultimately applied to ICANN to operate the .insurance and .bank top-level domains. Actually, you can go to the next slide, please.

What makes our TLDs among the most secure on the Internet today are several factors. First is the registrant verification that we do upfront and annually thereafter. The fact that these TLDs are highly restricted. And there's also a host of security protocols that we require registrants use to protect their domains.

Something that we think about broadly when we talk about security at fTLD is, is there a means to do something hierarchically? So that means at the TLD level versus at the traditional second-level domain or SLD level. A good example of this happened in 2018 when fTLD's TLDs were the first non-Google owned TLDs to be added to the Chrome HSTS preload list. Which in layman's terms means that domains are not reachable on the Internet, they're not properly secured with TLS certificate.

It was also in 2018 that we began considering a new approach to email authentication which is one of our most important security requirements. Specifically that domains in our zones must have a DMARC policy. DMARC stands for domain-based message authentication reporting and conformance. And for domains that are used for email, they must be at enforcement or reject within 90 days.

They also must have what's called a sender policy framework record or SPF. For those maybe less familiar with SPF, it's the list of authoritative

IP addresses that are permissible to have email sent on behalf of that specific domain.

For domains that are not used for email, we do require that DMARC be set at reject.

In addition to requiring email authentication records for our domains, another unique aspect about .bank and .insurance is that for domains to be in our zones they also have to be signed with DNSSEC and they also have to have other name servers in the .bank or .insurance zones. That requirement creates a different flavor of an NXDOMAIN.

If we fast forward to 2018, fTLD formed a working group, whose members you'll see listed on Slide 9 when we get to it, to explore the possibility of implementing DMARC at the TLD level, or that is above the organizational domain. The purpose of this is to add protection for nonexistent domains, to enforce compliance with our security requirements, and to provide DNS threat intelligence to mitigate DNS abuse. If you could go to the next slide.

For those of you who may be a little less familiar with email authentication and the protocols related to that—which are DMARC, SPF, and DKIM and they've been around for many years—this slide depicts what happens when a query is made in the DNS.

You can see at the bottom either a legitimate or an illegitimate email is sent. It goes out to the DNS. The DNS looks to see whether there are SPF, DKIM, and/or DMARC records. And then you can see it comes back down, and a few things can happen.

If the DMARC record exists, the message passes the DMARC test and the authenticated email goes to the recipient's inbox, which is a good thing.

If there's a DMARC record that exists but the DMARC fails—so perhaps the email is being sent from an SPF record that doesn't include the proper IP address—a few things can happen. You can see if the DMARC policy is set at "quarantine," that email is supposed to delivered to the recipient's spam box. If the record is set at "none," which is effectively nothing, the email goes directly to the inbox of the recipient. And if the policy is set at "reject," it would get trashed.

So that's what happens if a DMARC record exists. In the absence of a DMARC record, you can see that the receiver dictates the policy. What we understand historically is that 99.9% of the time it means the email goes into the recipient's inbox. And this is the big problem because recipients don't always know how to check for whether something is legitimate or illegitimate now.

To reiterate, the primary purpose of PSD DMARC is to protect nonexistent domains from being used to perpetrate DNS abuse and to ensure that fraudulent emails are not delivered or minimally put into the quarantine box of the recipient.

A really critical distinction to make here is that PSD DMARC will only come into play if there is not a DMARC record on the second-level domain. So when the email is sent and the query is made and the DNS is looking for DMARC records, it's going to look at the DMARC record first for the second-level domain. And if it finds it, it follows this path according to what the second-level domain policy is set. However, if

there's not a DMARC record for the second-level domain, it then would follow what has been set for PSD. And at least for .bank and .insurance that would be at "reject." Let's go on to the next slide.

Just a real quick example. This is our domain fTLD.com. You can see DMARC is set at "reject." You can see down below the SPF record includes all of the IP addresses that are authorized to send email on our behalf. Following again that chart on the prior page, if the DMARC policy checking passes, then we should be getting email in our fTLD.com email addresses. And if those DMARC policy checks fail, ideally we shouldn't be getting those emails at all. Next slide, please.

To get to the heart of PSD, our public suffix domain DMARC, when we talk about PSDs it could be something like .bank or .insurance. But it also includes TLDs where the organizational domain starts at a lower level. Some examples of this are the Canadian Centre for Cyber Security that operates .gc.ca and also the National Cyber Security Centre in the U.K. that operates .gov.uk.

PSD DMARC is a straightforward extension to DMARC whereby a determination is made about what the organizational domain is. PSD DMARC is ultimately only for a limited namespace such as when DMARC is required, such as what it is in .bank and .insurance.

It's also where registries have complete control over their second-level domain space. So a good example of that is our brand TLDs and then of course the examples I mentioned previously.

So ultimately there's going to need to be some type of registry or some type of means for the DNS to be able to identify whether the organizational domain is a PSD or not.

One of the key values of PSD and one that I mentioned a couple of times is that it adds an explicit policy for nonexistent domains. And again, it's usually "reject" and that certainly would be the case for .bank and .insurance. And given the feedback data element of DMARC, measures must be taken to minimize the potential for privacy considerations. Next slide, please.

Now that you understand the intention of DMARC, let me talk about the path forward. The good news is that RFC 9091 was published in July of this last year. It was after a lot of work done by fTLD, our working group, and of course the IETF. There is work underway now to update the current DMARC standard, which is RFC 7489, to include aspects of RFC 9091. fTLD and similarly situated gTLDs will need to seek permission from ICANN to implement this, and I'll get to more about that in a moment.

We do need to continue to work with mail service providers like Google and Microsoft and Yahoo because they have a role in the technical implementation of RFC 9091. We're fortunate that Microsoft has been a member of our working group since the very beginning. And now that the RFC has been published, our conversations with them have elevated to the fact that they are poised to be ready to support the RFC with a little bit of work for a specific number of TLDs.

We do continue to have stakeholder engagement with TLDs like ours that would have an interest in this. Brands, as I mentioned, probably are going to be the second most interested behind verified TLDs like ours.

Then lastly is to note that PSDs such as .gov, .uk, and .mil have already done this for their TLDs. They've put in DMARC records, and they've synthesized the role of in this case Microsoft or Google to validate the results. And they have reported no issues, which is really great news. Next slide.

So having talked about the path forward, I noticed I mentioned needing ICANN authorization to implement PSD DMARC. The reason for this is that for all of the new gTLDs that have been introduced since 2012, we have a specific prohibition in our registry agreement from adding a DMARC record into the DNS which is in effect a text record and that's what is in fact prohibited.

In contrast, it's not an issue for ccTLDs or non-ICANN gTLDs such as .gov and .mil because neither of those parties have contracts with ICANN. To continue the engagement with ICANN that's literally been happening probably for three years to brief them all along the way on the nature of our work. We've also briefed the SSAC on this as well.

In August of last year, we submitted an informal RSSAC request for this authorization to dive a little bit deeper into the consultation, and we are still waiting on feedback from ICANN to that informal request.

To be clear, PSD DMARC doesn't have technical risk because there are no root DNS changes. And fTLD's implementation for .bank and .insurance would mirror that, which has already been done by .gov, .uk, and also .mil. And as I mentioned earlier, there have been no technical issues reported. And I did also mention earlier that there are policy considerations for privacy issues. We believe that they'll largely be mitigated by constraining PSD to specific types of TLDs such as ours and that ultimately we'll need to get some guidance from ICANN once our RSEP has been approved. Next slide, please.

In summary, PSD DMARC extends coverage to organizational domains to mitigate DNS abuse and improve brand protection. The technical standard has already been defined by the IETF in RFC 9091. We already have good feedback from ccTLDs and non-ICANN gTLDs that have deployed it. The technical approach is low risk for expanding it to other gTLDs where it makes sense such as in our TLDs. And as I mentioned on the last slide, we'll need some guidance and feedback from ICANN on appropriate usage. Next slide, please.

A final bit of information. You can see the organizations that have been involved in our working group, and it's quite impressive, including governments domestically and abroad as well as the most well-known names in email authentication standards. There are quite a few PSDs interested in this type of approach to DMARC, and you can see a few of those listed here.

If you do ever want to look up the DMARC or SPF record of a second-level domain, you can use this easy tool. And then lastly, my contact

information should you have questions either now or at another time. So again, I thank you for your time and listening, and I'm happy to answer any questions you may have.

EBERHARD LISSE: Thank you very much. Again, all the addresses of all presenters or organizers of the roundtable are in the agenda that is published. So of course we will publish the PDFs of all the presentations, but if one wants to correspond with one of the presenters, just click on the corresponding links.

Now a question from Hesham: "Are there any guidelines for implementing it for ccTLDs?" Probably Craig is not the most ideal person to answer that because he works on a gTLD. ccTLDs can do whatever they want as long as they do it consistently. So it depends. You must ask each ccTLD manager themselves.

CRAIG SCHWARTZ: I'd also add here—and you're absolutely right about the authority of ccTLDs—but because mail service providers have a role in the implementation of PSD DMARC, there is something that they need to do for this to work. So for any ccTLD that wanted to try and do this now, I would suggest that they contact the folks at .gov.uk or maybe even .mil to see what they've done to synthesize the mail service provider role as that would enable the effective implementation absent mail service provider support at this time.

EBERHARD LISSE: He then [thought it up] from the technical perspective. I don't know an answer to that, but you answered it just now, I think.

CRAIG SCHWARTZ: I think so.

EBERHARD LISSE: Gordon Dick from Nominet: "More of a comment than a question which may help in use cases, and this may actually provide for wider coverage than just .brands. Open TLDs may wish to prohibit reserve names being spoofed, can use DMARC to do so. In .uk Nominet has utilized DMARC records against some reserve names to limit abuse. For example, .gov.co.uk cannot be registered and [thus] DMARC has successfully reduced spam pretending to be government.

CRAIG SCHWARTZ: Yeah, I think this is a great point and something I had not considered, but thank you for noting that.

EBERHARD LISSE: All right, I do see at the moment there is a hand from an attendee, Ulrich Wisser. You must ask your question in the Q&A pod, please. And the same for Nadia [inaudible]. We only take questions from the Q&A pod or if I see it in the chat. We have a few minutes left. I'm not going to push too hard, but if there are no further questions in the Q&A pod, you can always take it up offline by corresponding with Craig or any [related guest].

Okay, thank you very much. Interesting presentation. I learned a little bit from that, and I hope we can meet in person. Thank you.

CRAIG SCHWARTZ: Thank you.

EBERHARD LISSE: Okay, so next is Gustavo Lozano from ICANN. He is going to speak about the DANE interface to the MoSAPI project that they have. Can you please unmute yourself, uncover your camera, and either share the screen or Kim shares?

GUSTAVO LOZANO: Yeah.

EBERHARD LISSE: I can see you. I can hear you.

GUSTAVO LOZANO: Can you hear me?

EBERHARD LISSE: Yes, I can see you and hear you. So we're just waiting for the share screen or the slides from Kim.

GUSTAVO LOZANO: Yeah, I will share my screen. Just give me….

EBERHARD LISSE: Excellent.

GUSTAVO LOZANO: Can you see my screen now?

EBERHARD LISSE: Yes, we can. Go ahead. We have time, so don't have to rush.

GUSTAVO LOZANO: Okay, thank you. This is Gustavo Lozano from ICANN as Eberhard introduced me. I'm going to talk about the MoSAPI TLS client authentication. This is a project that we have been working for a while. And now that the IETF is working also on how to do TLS client authentication with them, I think this work may be of interest for ccTLDs and gTLDs alike.

I'm going to start with explaining what is the SLA monitoring system. This is in order to provide some overview or some context of how this MoSAPI works and how this will help your ccTLD.

When the new gTLD program was started, as part of the contract there was the idea to have ICANN monitoring the gTLDs. The idea was to be sure that there was compliance with the service level requirements and the service level agreement that was defined in the contract.

In order to do that we, ICANN, we have been working on this monitoring platform since 2012, and the first version was released in 2013. So what

is this monitoring platform that we call the SLAM or SLA monitoring system?

Well, it's Zabbix. Some of you may be familiar with this open source monitoring platform. We use exactly this monitoring platform plus some custom code. Zabbix, the company behind Zabbix the monitoring platform, they developed this custom code for ICANN in order to be able to monitor the gTLDs initially. And now we're monitoring basically all TLDs and [inaudible] using this platform. Besides Zabbix, we also have other parts of the code that have been developed internally.

Because this is a really interesting platform, according to the contract we have four hours to contact the TLD or the gTLD and try to resolve the issue. If the issue cannot be resolved and there is the potential possibility of using an emergency back-end operator in case that the registry operator is not able to backup their services online, then we have four hours basically to act. So this system needs to react in real time.

Basically what it does, when we detect an incident and there is an issue, the system automatically makes phone calls to the internal ICANN team and to the registry operator and registry service provider so they can start working on fixing this issue. So it's kind of complex, the code that we have now. And we have a lot of controls in place in order to mitigate the possibility of a false positive. Obviously, we don't want to call someone at 2:00 in the morning to say there is an incident going on and this may not be true. So that's why we have been investing all this time.

We have a probe node network. We have around 40 probe nodes, and we have all these probe nodes in all the ICANN regions. So we have nodes in America, Africa, Europe, all over the place.

We have centralized servers. These servers compile all the information from these probe nodes and they react. Basically, they make the phone calls and emails and provide some user interface so that our NOC can have a sense of how health the DNS is in real time.

We have a Network Operations Center. This is operating 24/7. And we have ICANN staff on call also 24/7 in case there is an issue so we can react promptly to any problem that may arise.

This is just a diagram of how it works. Basically, you have all these probe nodes around the world. They monitor the registry operator or the registry service provider, depending on the service model of the TLD or gTLD server. And then all the information is sent to the data processor that compiles all this information and determines if there is an incident or not.

Just to give you an idea, in order to declare an incident, we need to have 51% of the probe nodes that we [consider] online detecting the issue in order to declare an incident. On every probe node we have algorithms to detect if the IPv4 and the IPv6 connection can be considered healthy. If that connection to the Internet is not considered healthy, then we don't use the metrics from that probe node.

And all the probe nodes are located in basically data centers with carrier grade connectivity. We have multiple connections to the

Internet on every data center, and we try to have IPv6 native connectivity in all of them. We have in most of the cases IPv6 connectivity. Not in all of them, but we are working to have in all of them. I think that 90% of the probe nodes now have IPv6 connectivity.

So that's how the system works or how the monitoring system works. On your screen you can see the service level requirements that apply to gTLDs. These are the service level requirements that the system uses. So when we say that we monitor a ccTLD, we are monitoring the ccTLD from the perspective of these service level requirements.

As you can see, for example, for DNS TCP we have a minimum of 1,500 milliseconds to consider the query to be resolved. If not, we consider the query to be unanswered. For UDP it's less than 500 milliseconds from the perspective of every probe node. We also have service level requirements for RDDS which in the ICANN lingo means WHOIS or [web WHOIS]. And EPP, I think that everyone is familiar with the term EPP.

In addition to the service level requirements or what you call the service level agreement which is a typical service level agreement that you will find on a contract with an ISP, in the ICANN contract with the gTLDs we have what we call the emergency thresholds. These emergency thresholds are the thresholds that if they reach the 100%, then ICANN may—and I'm going to be very clear, may—trigger what we call the EBERO which is the emergency back-end registry operator. So in that case that there is an imminent failure of a gTLD and we decide that it's better for the registrants to have the emergency operator take over the gTLD, then that may be the decision.

And these are the thresholds that we have. So, for example, you can see that for DNS it's four hours. And these are computed on a rolling basis. So it means that we consider all the metrics from the past seven days, and on the next minute it's going to be the metrics of the past seven days minus one minute. So basically, the calculation is in real time and it [progresses as a] rolling week.

So those are the emergency thresholds. And as I was mentioning, you can see that we have really short period of time to react if there is an issue.

So after we started working on this SLA monitoring system there was the idea to provide all this information to gTLDs and also ccTLDs. So we developed what we call the monitoring system API. It's basically a RESTful API that all ccTLDs and all gTLDs have access to. Obviously, you need to get the access, but it's pretty simple.

If you use this API, you can get all the metrics that we capture from all the different probe nodes. You can also get access to the metrics when we declare an incident. How long that incident took to resolve and so on and so forth.

Also, by using this API you have access to the DAAR, domain abuse reporting monitoring from ICANN. So you if you get access to the MoSAPI, you get access to all this real time data from the SLA monitoring system plus you can get access to the DAAR. So basically, you have access to a lot of information that is really valuable for a ccTLD and a gTLD.

What are the advantages of this or the benefits of this API? Well, you have access to almost real time data. You have access to the continuous test data of the DNS. You have access to the DAAR statistics, as I was mentioning, for your TLD. And we have seen that some ccTLDs and some gTLDs also use this monitoring platform as a proactive monitoring tool. Meaning that they complement their own monitoring systems with the data from this SLA monitoring system. So if you have your own internal monitoring system, you can also feed from this system, compare your results, and see if your monitoring system matches what the SLA monitoring is seeing, for example. So those are the benefits that you can get if you get access to this API.

As I was mentioning, who can use this MoSAPI? Well, all gTLDs and all ccTLDs.

Now I'm going to go into the implementation of the TLS client authentication using [the SLA] records. Basically, it's what you would call using DANE but for TLS clients or to authenticate the clients, not the servers. That's basically what we have been doing for the past months, and this implementation has been in production since six months ago basically.

So what is the issue that we have in the gTLD space? We have, ICANN, we have contracts with several registry operators, thousands of them. And sometimes the registry operator is not the same entity as the registry service provider. So there are registry service providers that provide services to hundreds of registry operators, and obviously they need to have access to all this monitoring data.

Unfortunately, the way the business model in ICANN was developed was to have a contract with the registry operator. So the entity that is facing ICANN from the contractor perspective side is the registry operator and not the registry service provider. So that creates an issue that we're trying to solve with this implementation.

So let me go into the details. In the past, MoSAPI only offered HTTP basic authentication, so using a username and password. The credentials were managed by the registry operator, and that was really interesting because that means that the registry service operator needs to make these credentials available to the registry service provider. So this registry operator will go to our portal, will get access to find the username and the credentials and all of that, and then they will need to pass these credentials to the RSP (the registry service provider).

There was only one set of credentials per TLD. So if you were a registry service provider providing services to hundreds of TLDs, you will have hundreds of credentials, one credential per TLD. As a registry service provider you will need to establish multiple connections for all the different TLDs that you provide services to. So that was painful for the registry service provider.

They approached ICANN and said we have this issue. We need to manage all these hundreds of credentials. It's complex. Then sometimes the registry service operator, they don't know how to request the credentials. Then once they request it, there is not an easy way to securely provide this credential to us and so on and so forth.

So we started looking for a solution to this problem and after having conversations with the registry service providers, we noticed that TLS client authentication could be an option. At the end of the day in the DNS industry, because we are very familiar with EPP and EPP obviously works over TCP and most of the registries use TLS. So we came to the conclusion that using TLS client authentication could be an option because in the industry we are familiar with TLS client authentication at the end of the day.

So how it works. The registry operator, and this is the entity that has a contract with ICANN, once they go through our portal and they say I want to provide access to my registry service provider. So they need to provide a domain name in which the TLSA records will be found, and they also need to provide which roles they want to give access to the registry service provider.

Because to make things more complex, a registry operator may have multiple registry service providers for a TLD. So they may have a registry service provider for the DNS. They may have a different registry service provider for EPP and WHOIS and so on and so forth. Or they may have multiple providers for DNS.

So they can provide different access with different roles depending on the registry service provider. So the two key elements are the domain name where ICANN is going to retrieve the TLSA records and the roles that are authorized for that particular domain name.

What MoSAPI is going to do is retrieve all these different TLSA records from all the domain names that have been provided through the portal.

And then for the RSP it's pretty simple. They just go to the MoSAPI. They use TLS client authentication. They provide the certificate and if the certificate matches based on the TLSA record, then we provide access and we authorize based on the roles that were predefined in our portal.

You can have that domain name multiple times, so a particular registry service provider may use the same certificate for multiple TLDs over one connection if they decide to do so. So we also get less connections to our server because they can use one connection to access all these different TLDs.

This is an example. In this example, rsp1 is getting access to example01, example02. They are getting access to all the different roles based on the different domain name. Well, the domain name should be different here. That's a mistake on my part on that table, but the idea was to show that you can have different domain names like rsp2, rsp3 that make that example. But hopefully you get the idea.

What are the benefits of this model that we're implementing now in MoSAPI? Well, there is no sharing of credentials between the registry operator and the registry service provider. The only thing registered with ICANN is the domain name and the roles that will have access for that domain name. There is no need to manage passwords.

There is the ability to obtain data for multiple TLDs using one connection. Obviously, no need to get multiple credentials. You can use the same certificate to get access to all the TLDs that you want. And multiple parties can have access to the same role for a given TLD. So

you can have different registry service provider like the DNS provider, like the EPP provider getting access to all the information.

Once the registry has set this configuration once, only once, then the registry can manage their credentials. So they can change TLSA records so they can authenticate a new certificate if, for example, the private key was compromised. And there is no need to change anything because at the end of the day ICANN will just go to the domain name frequently. I think we have the timer every minute to get the TLSA record. So it's basically quasi real time. If you change the TLSA record, then you can have a different certificate being authenticated and getting access to the MoSAPI.

So that's how it works. Just some technical details. These are the following combinations of certificate usage types, selector and matching types that we support. We select these to make things easier. In the future if there is a need or someone believes that we should support different certificate usage types, we could do it. But for now these are the only ones that we are supporting. As I mentioned, we want to make things easier for now.

These are the public key algorithms that we are supporting on the client certificates. Basically, you can use RSA if the key size is 4096 bytes or higher, and you can use elliptic curve cryptography. And these are all the signature algorithms that are supported on the certificates. So basically, we use sha256 and up. And for elliptic curve it's exactly the same. We use SHA256 and above.

So I'm going to present the simple steps that a ccTLD or gTLD can do to test this implementation. It's pretty simple. The first comment will just create a new certificate. This is the certificate that we're going to use to authenticate with the MoSAPI.

Then after you create the certificate, you can use the DANE tool. This is a tool, I remember. I don't remember which implementation provides this DANE tool, but you can just search on the Internet. But you can use this DANE tool to generate the TLSA record for the certificate. So very simple. That's the TLSA record that you have there. What you are interested in is the [R] data. So basically the information after the "IN" for Internet. So "TLSA (03 01…)" all of that.

Then you just need to put this TLSA record into a domain name. For example, in this case we're using the domain name "tls-client-example.example.com." So basically, that's going to put that TLSA record on the public DNS.

And then you go to ICANN and say for this TLD please provide access to this domain name that you have there. And after ICANN has configured this, in the case of gTLDs that's through the portal, then it's pretty simple to get access. You just do something like a curl. You just set the certificate and the private key, and you just go to one of the endpoints and you will get access.

You are going to be authenticated using TLS client authentication using the TLSA record. And if you're authorized for that specific endpoint, then you get access to the data. And that's how it works.

If you want to request access, well if you're a gTLD, you just go to the portal. You already should have credentials to the portal, so we already have a secure way of authenticating you. And if you're a ccTLD, you request access. You send an email to GlobalSupport@icann.org, and then we'll use a process in which we go to the admin contacts and [technical] on IANA in order to authenticate that we are talking with a ccTLD operator. It's a pretty simple process, so that's the way you can get access.

And it's time for the Q&A, so I don't know who is going to control the Q&A.

EBERHARD LISSE: Thank you very much. I will control this as usual, and we are not in a hurry. So we can take the questions from the pod. There is one hand in the attendants list. I'm just saying for [inaudible], you must go into the Q&A. We will not take questions with the hand. Okay, anonymous attendee asks, "How MoSAPI provide monitoring of EPP service? If I understand correctly, to monitor EPP service MoSAPI has to act as an accredited registrar. But different registries have different accreditation requirements for registrars including technical requirements. So do you have some common policy with technical requirements for registries aimed at providing EPP monitoring service, or how do you solve the problem of different technical and [law] requirements of registries related with access to the registry database via EPP?"

GUSTAVO LOZANO: So right now we [are done] monitoring EPP. That is something that is in the works, but we have not been able to activate the EPP monitoring. One, when we developed the monitoring implementation back in 2013, the idea was exactly what you are describing. We will act as a registrar. We will executive some domain update comments, and then we will monitor how long it takes for that domain update to be propagated to the DNS. Basically, the idea was to include a timestamp in the name of the name server. So when we go to the public DNS and we get the name of the name server we will get the timestamp when we request the execution.

But right now we [are done] monitoring EPP. The system when we developed that, we included modules to support different extensions for the different implementations like [name store], etc. but we never activated. Obviously, to activate EPP it's a huge complex operational work. Our portal is designed to support the registry service providers to give us access through the username/password/certificate, whatever is needed to authenticate the EPP server. But we have not enabled that yet. So what we are monitoring right now is DNS and [inaudible]. In the ICANN lingo that means WHOIS and [web WHOIS]. And in the future it's the expectation that we would also in RDAP.

EBERHARD LISSE: Okay, Calvin Brown asks: "This is not registry component [inaudible] mentioned in some of the gTLD contracts to check on EPP stuff, right? How is that going?"

GUSTAVO LOZANO: Yes, it was mentioned that we are not monitoring EPP yet, and when we monitor EPP the idea is to actually have a connection to the EPP server and execute comments like domain update.

EBERHARD LISSE: And then Hugo Salgado from .cl: "Can you have multiple TLSA records with different certificates for the same domain, or do you need to define multiple domains?"

GUSTAVO LOZANO: I'm pretty sure that you can have multiple TLSA records on an owner name [inaudible] we will get all of them. I'm 99% sure, but I can confirm that with the technical folks and be sure that we support that. For sure what I have tested several times is you can have multiple domain names for the same TLD and that works. But having multiple TLSA records on one owner name, I'm pretty sure that it works but I will check and come back to you, Hugo.

EBERHARD LISSE: Thank you. And I don't see any more questions.

GUSTAVO LOZANO: I see a lot of hands or a few hands.

EBERHARD LISSE: Yeah, but sorry. The panelists can speak. The attendees must please put it into the question and answer pod.

GUSTAVO LOZANO: Oh, got it.

EBERHARD LISSE: Under panelists, Stephen Deerhake has a question. Stephen, you have the floor.

STEPHEN DEERHAKE: Thank you, sir. How about an easy to access web-based dashboard so that I as a registry operator can see that I'm "in compliance"? Thank you.

GUSTAVO LOZANO: That's a good idea. I believe that in the past there was this proposal at least for gTLDs, which they have access to our web portal, to have some kind of dashboard in which they can see compliance with different aspects, technical aspects. I don't think that is available yet, but it's something that has been considered in the past. But again, this is for gTLDs. That's what we have considered in the past as [inaudible].

EBERHARD LISSE: I only see the hand in the attendees which we are not going to take because the can type into questions and answers. No more questions then. If this person, "AAAA," needs to have a question, they must take it up with you directly.

GUSTAVO LOZANO: Yeah, sure.

EBERHARD LISSE: [inaudible]

KATHY SCHNITT: Eberhard, it does look like he put his note in the chat.

EBERHARD LISSE: Oh, sorry. “What is the effect in processing and effectiveness in query response time?”

GUSTAVO LOZANO: Could you elaborate on that question? I don’t understand the question, sorry.

EBERHARD LISSE: “What is the effect in processing and effectiveness in query response time?” If you can’t understand it, then that individual should email you directly. The email address is on the agenda.

GUSTAVO LOZANO: Yeah.

EBERHARD LISSE: They can just email you, yeah?

GUSTAVO LOZANO: Yes, that's fine.

EBERHARD LISSE: Thank you very much. Okay, we are a little bit ahead of time. So if there are no more questions, we go and thank you very much for your presentation.

GUSTAVO LOZANO: Thank you.

EBERHARD LISSE: We had the same thing before and we will have it again when there are further developments. I like research and statistics. Thank you. Next presenter will be Dr. Rodriguez from .pr. I think you are unmuted. Can you just say a word?

PABLO RODRIGUEZ: Thank you very much.

EBERHARD LISSE: Yes, we can hear you. Thank you. Hang on. I just wanted to check. Thank you very much for Gustavo. And then as newcomers perhaps don't know but older participants know, we always allow or ask the host of the meeting even on a virtual meeting to give us a host presentation. In other words, a presentation on a topic of their choice. And Dr. Rodriguez, you have the floor.

PABLO RODRIGUEZ: Thank you, [inaudible]. And thank you all for giving me the opportunity to share an update several things that are taking place in Puerto Rico. For example—next slide, please—for those of you who graciously came to participate at ICANN61 in 2018, as you know we had just passed a terrible Category 5 hurricane that devastated the entire island. And as you can see in the picture that I'm sharing, to the right in your screen you will see that there is a plot of land that at the time was being constructed and some construction activity was taking place. Even though most of the island had been destroyed, we continued marching forward.

That was then. I'd like to show you what it looks like now. This is what it looks like now. A tremendous entertainment center known locally as Distrito T-Mobile or the T-Mobile District, entertainment district, has a tremendous amount of restaurants and bars and so on.

So let me show you what can you expect. For example, we have a number of new hotels in that area. In addition to that, we also have—and this is right next to the convention center. This is literally crossing the street and you're in it. There are new hotels, and the entertainment center contains ample spaces where we can all meet and share discussions and so on. There is a concert hall. There are a number of movie theaters. There is also a huge entertainment place for children of all ages. And you have all sorts of video games and activities and so on.

And most recently they also have this [sling] that goes from the Distrito T-Mobile straight into the convention center. So that is another way in

which you can reach the convention center there. And there are a number of other areas where you can meet and gather—pubs, restaurants, pizzerias, and so on. So if you would like to find out more about what is waiting for you next time you'll be in Puerto Rico, you can visit this URL and you will find further information.

But also, I'd like to share with you what we are up to in [inaudible] .pr. I'd like to introduce you to tu.pr which means your.pr. This is intended exclusively for residents in Puerto Rico. Necessarily you don't need to be a citizen of the United States, but as long as you're residing in Puerto Rico you are entitled to the benefits of this package.

What we're trying to do is to provide individuals and small businesses with an opportunity to have a presence on the Internet. As you all know and as we have shared through the years with the pandemic we have found that many people finally realize as a consequence of the pandemic that it was important to have a presence on the Internet. And many of them were caught unprepared to have that presence.

To this day there are many people that understand that they need to have a presence online but don't know how. They don't have the skills. So what do I do to help to facilitate a presence on the Internet for the baker, for the laundry owner, for the grocery store owner, for the small and medium size type of businesses? What can we do? So what we attempted to do as you will see now is that we developed a new platform that is strictly designed for Puerto Rico. Next slide, please.

And it has a special price. As you may have seen on the previous slide, we're charging $99.99 strictly to Puerto Rican residents, to residents of

Puerto Rico. And of course, you will have to provide some evidence, light or gas bills, telephone bills, cable bills, that type of thing, or evidence that you are a registered merchant in Puerto Rico. And you can have access to this platform and register a domain name for $99.99.

And you will have a domain dashboard. You will be provided a domain dashboard. All of this is in cPanel so you will see much of this is very familiar to what you already know. And with this dashboard you will have the various features that normally you have on a dashboard which is what is the domain name, when was it registered, when does it expire. And you have the ability to manage your account and manage other services, get metrics.

And also, we provide a number of services that we understand are helpful for someone who is starting to have a presence on the Internet. And we are bundling all of these services and features into this package, so it should be enough for someone, for a neophyte, for someone who is beginning to have a presence on the Internet without much difficulty.

For example, we offer free DNS hosting. There you can see that in other many, many ISPs in many other companies this is an additional charge. In our case, we're including this with every domain registration. It's a free DNS hosting service that is accessible by the configuration tab in the domain dashboard.

So again, this is one way which we are using to help and facilitate services to our local customers. We are also providing a hosting platform. And this hosting platform is provisioned with Linux, Apache,

PHP, and MySQL. So it's a platform that is based on Linux, Apache, PHP, and MySQL.

Again, you can see how you can manage the various services that you would normally do on a hosting platform. For example, you can manage the capacities. And also as you grow, this service also grows with you. You can also choose to upgrade to additional services.

We are also providing the various [full] hosting features. For example, [full] hosting service includes a minimal set of tools to aid in the use of common domain functions. In other words, you can change your DNS. You can change your MX records. You can edit A records and so on. And so you have total autonomy to manage your domain name.

One of those features that seems to be very important to many people is the DNS zone editor. With that, we have provided the ability for individuals to manage their DNS zone records in this portal. And that provides a tremendous amount of freedom and autonomy to the end user.

Also, we have provided DNSSEC. And by that I mean that this service also allows our clients to sign their DNS zones with DNSSEC. The idea is that we want to streamline, facilitate our customers with all the security measures that we can possibly provide them that they necessarily won't understand or it would be difficult for them to adopt or to implement and so on. So what we have done is that we are bundling all of these services, once again, in this one platform so that you can have as many bases covered as possible. Furthermore, we also have a file manager.

EBERHARD LISSE: You're on mute.

PABLO RODRIGUEZ: Can you hear me now?

EBERHARD LISSE: Yes, it's better.

PABLO RODRIGUEZ: Thank you. Apologies for that. I don't know what happened there. So we also have a file manager that we provide our customers. And aside from the FTP access, the file manager is accessible from the hosting panel to publish a simple website that should not and cannot exceed a total of 100 megabytes.

But we also add email accounts, and these email accounts also cannot exceed the 100 megabyte storage quota. But the idea is that you have the minimum bare bones that you need to have a presence online, communicating online in a secure fashion. And then if you need further, if you need more and as you grow you may need bigger services, then you can upgrade those services.

Finally, if you're one of those individuals that already has a website and you're only interested in having that domain name in the .pr universe, we provide domain redirection. That domain redirection allows individuals to point their domain names to their particular pages which

are already well-known by their existing customers and potential customers in other areas. Next slide, please.

I want to thank you all very much for this opportunity. And I'm extremely grateful not only for this opportunity. Thank you, Dr. Lisse, to the entire staff, and to all of you involved in this Tech Day that make this Tech Day possible. It is extremely valuable for our community. For a global community it is extremely valuable for all of us. And I want to thank you all for sharing with us when we needed it in 2018 right after the hurricane, and thank you for your continued support. Any questions, I'll be happy to entertain.

EBERHARD LISSE: Thank you very much. Nice presentation of what you guys are up to. Calvin Brown asked a little bit facetiously: "I'm interested in what the ISPs think of this. Also, what do the competition authorities think of this?"

PABLO RODRIGUEZ: Absolutely. We have a special revenue share program that we have devised, that we have designed for potential resellers. And I hope that ISPs will be glad that we have developed this program because it is a lower price than we normally charge internationally and it would allow them to enrich their offering of products and services to their current offering. Consequently, we expect that there will be an increase in domain registrations and we will be able to either provide services this way or we can [sign] agreements. As I already mentioned, we already

have an agreement we sign for resellers and these ISPs can [inaudible] big part in this special price program for Puerto Rico residents. And they can participate in that, and we expect them to be very, very successful. I hope I answered the question. Did I miss anything?

EBERHARD LISSE: In Namibia we also have, or in .na we also have got different price so to speak for Namibians or permanent residents or registered companies and foreigners. And our competition commission is only interested in Namibian clients or consumer protection of Namibian clients. What foreigners do is not an issue of the local competition commission. I assume this will be similar elsewhere because nobody forces anyone to register a domain in a different ccTLD.

PABLO RODRIGUEZ: Our experience is that major companies in Puerto Rico interestingly, they go and purchase services from registrars in the main U.S. specifically because I believe that they have other domain names that they're handling for them so .pr will also be added into the cache of services that they are offering. However, with this program now we can aspire to obtain this additional price and we are looking forward to find how the market is going to react. We expect that it's going to be very positive.

EBERHARD LISSE: I don't see any further questions. And since that is the last one before the break, everybody is released for the break. Afterwards, we will have

the presentation from Kathleen Moriarty and Paul Vixie, then the roundtable, and then the .ua presentation. So all be back.

PABLO RODRIGUEZ: Thank you, Dr. Lisse. Thank you all.

[END OF TRANSCRIPTION]